



10/520945
PCT/GB 2003/003051
Re PCT/PTO 12 JAN 2005
INVESTOR IN PEOPLE

PRIORITY DOCUMENT

SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)

The Patent Office
Concept House
Cardiff Road

Newport
South Wales
NP10 8QQ

FRECD 09 SEP 2003

PCT

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

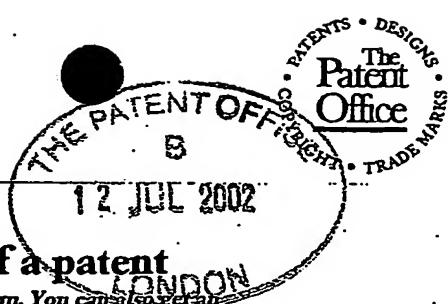
In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 20 August 2003

15 JUL 02 E733 07-3 D 2000
P01/7700 0.00 0216257.6

1/77

Request for grant of a patent

(See the notes on the back of this form. You can also obtain an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office

 Cardiff Road
Newport
South Wales
NP10 8QQ

1. Your reference

PDG/AJB/24719

2. Patent application number

(The Patent Office will fill in this part)

171511 0753

0216257.6

3. Full name, address and postcode of the or of each applicant (*underline all surnames*)
 Snell & Wilcox Limited
6 Old Lodge Place
St Margaret's
Twickenham
Middlesex TW1 1RQ
Patents ADP number (*if you know it*)

55797 84003

If the applicant is a corporate body, give the country/state of its incorporation

United Kingdom

4. Title of the invention

Improved Noise Reduction.

5. Name of your agent (*if you have one*)
 "Address for service" in the United Kingdom
to which all correspondence should be sent
(including the postcode)

 MATHYS & SQUIRE
100 Gray's Inn Road
London WC1X 8AL
United Kingdom
Patents ADP number (*if you know it*)

1081001

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (*if you know it*) the or each application number

Country

Priority application number

Date of filing

*(if you know it)**(day / month / year)*

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
*(day / month / year)*8. Is a statement of inventorship and of right to grant of a patent required in support of this request? *(Answer 'Yes' if:*

YES

- a) any applicant named in part 3 is not an inventor, or
 - b) there is an inventor who is not named as an applicant, or
 - c) any named applicants is a corporate body.
- See note (d))*

Enter the number of sheets for any of the following items you are filing with this form.
Do not count copies of the same document

Continuation sheets of this form

Description 12

Claim(s) 2

Abstract -

Drawing(s) 7 + 2 for

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

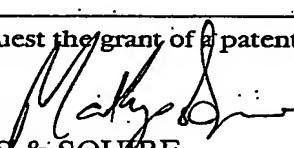
Request for preliminary examination and search (Patents Form 9/77)

1

Request for substantive examination
(Patents Form 10/77)Any other documents
(Please specify)

11.

I/We request the grant of a patent on the basis of this application

Signature 

Date

MATHYS & SQUIRE

12 July 2002

12. Name and daytime telephone number of person to contact in the United Kingdom

Peter D Garratt - 020 7830 0000

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

- 1 -

IMPROVED NOISE REDUCTION

This invention concerns the processing of video signals to change their subjective quality or to make them more amenable to other processes such as compression.

The reduction of noise in video signals, either by reducing the spatial 5 bandwidth, or by averaging stationary objects over several different temporal samples, is well known. These processes have been found not only to enhance the subjective appearance of pictures but also to facilitate the data compression of video by removing less-visually-significant information prior to the compression process.

10 A commonly-used technique is to combine the video signal from a point in a current frame with the signal from a corresponding point in a previous frame or frames. This is usually implemented as a recursive (IIR) filter where the recursion coefficient is varied in response to a motion dependant control signal. Stationary areas may be averaged over several frames, whereas moving areas may not be 15 averaged at all.

A known implementation of this technique is to subtract two consecutive frames to obtain a signal containing noise and motion; this signal is passed through a non-linear function which attenuates small signals but allows large signals to pass. The output of the non-linear processing is then added to the 20 earlier frame to obtain a noise-reduced output. In this arrangement the instantaneous gain of the non linear process effectively controls the degree of recursion. When the instantaneous gain is zero the previous frame is output; and when the gain is unity the added previous-frame information is cancelled by the subtracted previous-frame information, and the processing has no effect. Usually 25 the shape of the non-linearity, and hence the gain characteristic, is varied in response to a motion adaptation signal so that minimal processing is applied in moving areas.

In spite of this adaptation, the performance may not be adequate in moving areas and a further adaptation to a non-recursive, spatial bandwidth 30 reduction noise reducer may be used. Typically this spatial processing "cores"

the high-frequency content of the signal by attenuating low amplitude components (which are mainly noise), whilst allowing larger high frequency components to pass.

Where video images have been derived from film, the grain of the film may 5 be a subjective impairment and these known video noise reduction techniques may be used to make it less objectionable. However, film grain is a contribution to the well-recognised and appreciated "film-look" and removing it by these prior art methods sometimes leads to subjectively unacceptable results.

The inventors have devised a novel and flexible method of video 10 processing which can be used to modify video material to achieve one or more of the following:

- Reduction of video noise
- Improvement in the subjective appearance of film grain
- Compression pre-processing
- 15 · Alteration of the spatial frequency response of moving pictures to provide equalisation of distortions or subjective enhancement.

The invention consists, in one aspect, of a recursive video process in which spatially co-sited information from different temporal samples is combined, 20 and the combination varies in dependence upon spatial frequency.

Suitably, recursive and non-recursive filtering is combined.

In a further aspect, one or more non-recursively separated detail components are combined with the video signal.

Advantageously, one or more noise components are removed from one or 25 more non-recursively separated detail components, and the resulting noise-reduced detail is combined with the video signal.

In a yet further aspect of the invention, one or more noise components are separated from a video signal by a recursive process where the recursion depends on spatial frequency and one or more of the separated noise 30 components is combined with a noise-reduced video signal.

An example of the invention will now be described with reference to the drawings in which:

5 Figure 1 shows a video processing system in accordance with the invention;

Figure 2 shows an arrangement of spatial frequency bands;

10 Figure 3 shows a non-linear transfer characteristic;

Figure 4 shows an noise and detail processor;

Figure 5a shows a band splitting filter;

15 Figure 5b shows an alternative structure for the filter of Figure 5a;

Figure 6 shows a vertical band splitting filter;

20 Figure 7 shows a method of splitting the frequency bands of Figure 2; and

Figure 8 shows an alternative method of splitting the frequency bands of Figure 2.

25 Referring to Figure 1, a video input (1) feeds a one-frame compensating delay (2), the delayed output feeds a subtractor (3), whose inverting input is driven by noise which has been extracted from the signal by methods which will be described below. The noise is co-timed with the video so that the output of the subtractor (3) is a noise-reduced signal. This is combined in an adder (4) with an image enhancement signal (also to be described below) and forms the video
30 output from the process (5).

The input signal (1) also feeds a bank of spatial filters (6a), (6b) etc. These may be band-pass or high-pass, and are chosen to decompose the higher frequency part of the signal spectrum into a number of separate frequency bands. A typical arrangement of frequency bands is shown in Figure 2. This 5 figure shows nine two-dimensional frequency bands (21), (22) ... (29) which can be filtered from a video signal having 576 active lines and 720 horizontal samples per active line. (Note that nine filters (6) in Figure 1 are required to implement this arrangement; the low-frequency band (20) does not have a corresponding filter (6) in Figure 1.) The frequency bands should tessellate and the filter 10 characteristics should be "complementary" so that if the filtered outputs, and the low frequency band (20), are added together the original signal is reconstructed, substantially without distortion. Examples of suitable filtering methods will be described below.

If the video is interlaced it will be necessary to use lines from both fields in 15 order to separate the higher vertical-frequency bands (21) and (22). This is undesirable unless the signal is derived from film (or some other process in which the temporal sampling rate is nearer the picture rate than the field rate); for true interlaced video the bands (21) and (22) would not normally be used. Each of the filters (6) in Figure 1 (there will be nine filters to implement the 20 arrangement shown in Figure 2) feeds a respective "inverse-coring circuit" (7a), (7b) etc. These have a transfer characteristic of the form shown in Figure 3. Referring to this Figure, the characteristic has a linear region (31) where small input signals are passed transparently, and low-gain regions (32) and (33) where larger-amplitude positive and negative inputs are attenuated.

25 Returning to Figure 1, the outputs of the non-linear circuits (7) are fed to a linear combiner (8). This forms two different weighted combinations of its input signals; one combination (9) is optimised to separate the noise from the input signal, and the second combination (10) is optimised to separate high-frequency detail from the input signal.

30 The propagation delay of the filters (6) in Figure 1 is made equal to one frame (this will usually require additional delay elements, not shown in the

Figure) so that detail and noise signals (9) and (10) are co-timed with the output from the compensating delay (2). The noise can then be subtracted (3) to give a noise-reduced signal (11), and the detail can be added (4) to give a noise-reduced and detail-enhanced video output (5).

5 The processing described so far is purely spatial in character and it is known that spatial noise-reduction can cause undesirable artefacts when a high degree of noise reduction is applied. The system described above minimises these artefacts by applying different levels of noise reduction in the different frequency bands, but further improvement is possible by making use of temporal recursion, as will now be described.

10 The noise-reduced signal (11) is subtracted (12) from the input signal (1). These signals are separated in time by one frame, because of the compensating delay (2), and so, in the absence of movement, the difference comprises only noise. The difference signal is passed to a second set of band splitting filters 15. (13a), (13b) etc., and (14). These filters separate the signal into a number of bands in a similar way to the filters (6); however, a low pass filter (14) is included in this case.

20 The separated bands, including the low-frequency band, are processed in an array of inverse-coring circuits (15a), (15b), etc., and (16) which operate in the same way as the blocks (7) to attenuate all except the small signal components. The filtered, inverse-cored signals are fed to the combiner (8) where they are linearly combined with the signals from the blocks (7) to create the separated-noise signal (9).

25 The propagation delay of the filters (15) and (16) is arranged to be one frame (by adding compensating delays, not shown in the Figure, if necessary) so that the filtered contribution from the input signal (1) (via the subtractor (12)) arrives at the output of the combiner (8) co-timed with the output of the compensating delay (2).

30 By selecting appropriate weighting factors in the combiner (8) to generate the noise signal (9) both recursive and non-recursive noise reduction of the video signal can be achieved. Generally when both the recursive and non recursive

contributions from the same frequency band are used care is necessary to ensure that the corresponding signal weighting factors sum to unity, so that the band in question is not unduly emphasised or de-emphasised. The optimisation of the contributions will depend on the nature of the wanted signal and the noise;

- 5 as well as adjusting the combiner (8) it may be advantageous to alter the size of the linear regions of each of the blocks (7), (15) and (16) depending on the level of noise in the respective frequency bands.

The optimisation of the recursive contributions to the noise signal (9) should also take account of "shot-changes" and the like. If there were no
10 contributions from the blocks (7), one frame after a discontinuity there will be similar levels of noise at both input to the subtractor (12); optimum noise reduction is then achieved when the gain of the route from the output of the subtractor (12) to the negating input of the subtractor (3) is one half. This will give a 3 dB noise improvement because the current field is effectively averaged with
15 the previous field (as far as small signal components are concerned). One frame later, this noise improvement will be apparent on the signal (11) fed back to the subtractor (12), and its output will better correspond to the noise at the output of the delay (2); the corresponding optimum gain is then two thirds and the noise improvement is then 5 dB. For subsequent frames the sequence of optimum gain
20 values is ... 3/4, 4/5, etc. This is analogous to the process described in European Patent EP 0 893 024 B1.

In practice better noise reduction than this will be achieved because the contributions from the blocks (7) to the noise signal (9) will ensure that there is less noise on the feedback signal (11), and so a higher degree of recursion is
25 possible with consequent improved noise reduction.

The capabilities of the invention may be extended in another embodiment in which the operation of the block (8) in Figure 1 differs from that described above; Figure 4 shows this alternative arrangement of the block (8).

Figure 4 has two sets of input signals. A first set of nine inputs (41) are
30 inverse-cored sub-bands separated from the video input (i.e. the outputs of the blocks (7) in Figure 1); a second set of ten inputs (42) are inverse-cored sub-

bands separated from the difference between the input video and the noise-reduced video of the preceding frame (i.e. the outputs from the blocks (15) and (16) in Figure 1). Typically the bands would be as shown in Figure 2; although other arrangements and numbers of bands are not ruled out. Nine of the input
5 bands (42) cover the same respective spectral areas as the input bands (41); the tenth of the inputs (42) covers the low frequencies not included in the other bands.

A gain adjustment block (43) individually adjusts the gains of each of the 19 inputs (41) and (42), and sums together corresponding bands from the two
10 groups of inputs. The gain values are chosen so that when the ten outputs from the block (43) are summed in the block (44), an output signal (45) is obtained which corresponds as closely as possible to the noise content of the video signal being processed. (This is the signal (9) of Figure 1.)

The nine high-frequency bands from the gain adjustment block (43), i.e. all
15 except the low-frequency band, are respectively subtracted (46) from the inputs (41) to give a set of nine noise-free sub-band components which are input to a gain adjustment and summing block (47). The individual gains of the bands are chosen so that the summed signal provides a detail correction signal which will be added to the video being processed so as to achieve subjective
20 enhancement. Usually constructive edge enhancement is required but negative gains in the block (47) could be used to de-enhance edges. Because the enhancement signal has had the noise subtracted from it, a greater degree of enhancement is possible without creating undesirable artefacts.

There are applications when the objective is not to remove all noise from a
25 signal but to alter the characteristics of the noise – perhaps to match the noise characteristics of another signal. The outputs of the block (43) comprise a close approximation to the noise content of the input signal; fortunately this noise is presented as ten separate frequency bands, and it is therefore possible to modify the character of the noise by adjusting the gains of the individual bands before
30 adding them together. This gain adjustment can be done in the block (49), and

the resulting wanted noise can be added (410) to the detail correction signal (48) to give a combined noise and detail correction signal (411).

The output signals (45) and (411) from Figure 4 can be used as the signals (9) and (10) respectively in Figure 1. The detail correction and the wanted noise are not included in the recursion loop which makes them easier to optimise. However, the best possible approximation to the noise is subtracted from the video fed back into the recursion loop, thus allowing the majority of the noise on the input video to be separated and cancelled. This allows very flexible video processing where detail enhancement or de-enhancement can be adjusted independently of noise optimisation.

Clearly the effectiveness of the processing depends on the ability of the filter bands to discriminate wanted signal from noise, so that higher degrees of noise reduction can be applied in bands which contain little or no wanted signals. The arrangement shown in Figure 2 has been found to be effective and a method of implementing the necessary band splitting filters will now be described. For the purpose of this explanation an interlaced signal having 576 active lines, each having 720 horizontal samples, will be assumed; the skilled person would have no difficulty in adapting the methods described to other video formats.

As can be seen from Figure 2, the frequency bands can be obtained by a hierarchical sequence of band splitting operations. First there is a vertical split at 144 cycles per active picture height (c/aph); the resulting high and low vertical frequencies are then split by horizontal filters at 180 c/apw. The lowest frequency band (nearest the origin) from this operation is then split horizontally and vertically at 90 c/apw and 72 c/aph respectively. The lowest band from this second split is then divided at 45 c/paw and 36 c/apw to derive the remaining bands. Each of these band splitting operations (both horizontal and vertical) can be accomplished by using a symmetrical, three tap transversal filter of the form shown in Figure 5.

Either the form of Figure 5a or Figure 5b may be used; they are equivalents. This type of filter gives a cosinusoidal frequency response and introduces no group delay distortion. The horizontal band-split at 180 c/apw

shown in Figure 3 can be obtained with a delay of one sample, the split at 90 c/apw can be achieved with a delay of two samples, and the split at 45 c/apw requires a delay of four samples.

Because the filter response is periodic, it is necessary to consider the

5 higher-frequency repeats of the response. This is not a problem for the one-sample-delay filter as the higher order responses are above the Nyquist limit; however, in the case of the two-sample-delay filter, the effects of the higher order must be considered. Fortunately if the two filters are cascaded (with the wider band filter first) the higher-order responses of the second filter fall in the stop-
10 band of the first filter.

The vertical band-split at 144 c/aph requires a delay corresponding to one line pitch. Usually the signal is interlaced and so a signal delay of about one field is required. A suitable arrangement is shown in Figure 6. The video input (61) feeds a 313-line delay (62) and a changeover switch (63). When the switch is in
15 the position marked A the input also feeds a one-line delay (64). The output from the 313-line delay (62) feeds a weighted-input adder (65), which combines the 313-line-delayed input (weight $\frac{1}{2}$) with the undelayed input (weight $\frac{1}{4}$) and the output of the one-line delay (64) (weight $\frac{1}{4}$). The resulting output (66) forms the vertically low-pass filtered signal and corresponds to a combination of three
20 vertically adjacent picture lines, two from the input field, and the intervening line from the previous field. This low pass output is subtracted (67) from the central line to give a high pass output (68).

The combination of lines from adjacent fields (as described above) usually gives temporal as well as spatial filtering; however if the input is derived from film
25 (or some other source with a temporal sampling rate lower than the video field rate) it is possible to obtain purely spatial filtering by only combining lines from fields which correspond to the same point in time. In the case of film having a frame rate corresponding to the video frame rate, this is achieved by operating the switch (63) at field rate so that, when the input field (61) does not correspond
30 to the same point in time as the output of the 313-line delay (62), the switch (63) is set to position B so that two lines from an earlier field are used. These lines are

made available by the combination of the 312-line delay (69) and the one-line delay (64). (This method of processing is described in European patent EP 0 626 120.)

If the well-known 3:2 pulldown technique is used to synchronise a film frame rate lower than the video frame rate to the video frame rate, the switch (63) can be operated in a suitable pattern (such as A B B A B A B B etc.) to ensure that only time-coincident lines are combined.

The complete set of bands shown in Figure 2 can be obtained from a cascaded arrangement of vertical and horizontal spatial frequency splitting filters as shown in Figure 7. Each of the filters sums three, differently delayed versions of its input, weighted in the proportions $\frac{1}{4} \frac{1}{2} \frac{1}{4}$ as described above (see Figures 5 and 6). The relevant delay values (between the contributions which are combined in the respective filters) are shown in Figure 7, in sample periods for the horizontal filters, and picture lines for the vertical filters. As mentioned previously, the propagation delays of all the filters must be equal to one frame and this will require compensating delays to be inserted at appropriate points; these are not shown in the Figures.

The first vertical filter (71) uses information from two fields and can use the arrangement shown in Figure 6. However the vertical filter (73) requires a delay of two picture lines which equates to a delay of one line period, and the vertical filter (72) requires a delay of two line periods. These latter two filters can therefore use the structures shown in Figure 5. All the horizontal filters can use the arrangement shown in Figure 5.

As mentioned previously, compensating delays must be added to the various filtered signals so as to make the delay for each band equal to one frame. It is therefore advantageous to re-arrange the filters where possible so as to make their vertical delays equal. The processing below the line (72) in Figure 7 uses only linear combinations of delayed signals from one field; it is therefore possible to reorder these processes, provided that the contributions to each output remain unchanged.

Figure 8 shows a rearrangement of Figure 7 in which the bands 23 to 29 and the low-frequency band 20 are derived from a single set of six cascaded line delays. All the horizontal filters in Figure 8 have the same structures those in Figure 7. The vertical filter at 144 c/aph (81) is also the same as in Figure 7; the
5 two other vertical filters are differently structured, but achieve the same frequency responses for the various bands.

Referring to Figure 8, the processing above the line (80) is identical to that shown in Figure 7. The signal (83) has been vertically low pass filtered (81) (at 144 c/aph) so as to remove the bands 21 and 22. Two further vertical band
10 splitting filters are required, one (at 72 c/aph) to separate the bands 24 and 25 (from the bands 26, 27, 28, 29 and 20), and a second filter (at 36 c/aph) to separate the bands 27 and 28 (from the bands 29 and 20).

The signal (83) is fed to six one-line-period delay lines giving delayed outputs (a) to (f) inclusive. The band 23 is obtained by high pass filtering (81) the
15 output (c) of the third line delay.

The first vertically low-pass filtered signal is obtained by summing (84) the three-line-delayed signal (c) with the signals (b) and (d), which correspond to the succeeding and preceding lines respectively. The signals are weighted $\frac{1}{4}$, $\frac{1}{2}$, $\frac{1}{4}$ as in Figure 5. The resulting signal (85) is subtracted (86) from the central signal
20 (c) to give a high-pass output (87) which contains bands 24, 25 and part of band 23. This is low-pass filtered (88) to leave only bands 24 and 25, which are separated (as previously described) in the horizontal filter (89).

The vertical low-pass signal (85) contains the band 26 together with both higher and lower horizontal frequency components, which are removed by the
25 low pass filter (810) and the high-pass filter (811).

The second vertical band split must have a response equivalent to the cascading of the filters (71) and (73) in Figure 7. This is obtained by taking an appropriate weighted sum of the signal (83) and its delayed versions (a) to (f) in the adder (812). The resulting low-pass signal (813) is vertically co-timed with the
30 signal (c) and contains the bands 20 and 29 together with parts of the bands 26

and 23. These latter parts are removed by the low-pass filters (814) and (815).
The two wanted bands are then separated by the horizontal splitting filter (816).

The low-pass output (813) of the second vertical filter is subtracted (817)
from the low-pass output (85) of the first vertical filter to give the signal (818),
5 which corresponds to the bands 27 and 28, together with parts of bands 26 and
23. These latter unwanted high-horizontal-frequency components are removed
by two low-pass filters (819) and (820) before separation of the wanted bands by
the horizontal splitter (821).

It can be seen that outputs of all the filters below the line (80) are arranged
10 to be timed three field-lines after the signal (83) and so only horizontal
compensating delays are required to synchronise their output bands. A single
vertical compensating delay at the low-pass output of the filter (81) can be used
to synchronise these bands with the bands 21 and 22.

It must be noted that the above descriptions are examples and many
15 variations of the concept are possible. For example the number and shape of the
frequency bands may differ from those shown in Figure 3 and the invention is not
limited to interlaced formats or to formats having 576 active lines and 720
samples per active line. Transversal filters with different numbers of signal
contributions and weights from those shown may be used; the horizontal filters
20 could be implemented by inductors and capacitors. Not all the high frequency
bands in the recursive processing may correspond to bands in the non-recursive
processing, and vice-versa. The inverse coring circuits may differ from each
other, and their individual characteristics may be varied so as to optimise the
performance for a particular input.

CLAIMS

1. A video process in which spatially co-sited information from different temporal samples is recursively combined, wherein the combination varies in dependence upon spatial frequency.
2. A process according to Claim 1 in which recursive and non-recursive filtering is combined.
3. A process according to Claim 2 in which one or more non-recursively separated detail components is combined with the video signal.
4. A process according to Claim 3 in which one or more noise components are removed from one or more non-recursively separated detail components and the resulting noise-reduced detail is combined with the video signal.
5. A process according to any of the preceding claims in which one or more detail components are separated from a video signal and combined with a noise reduced video signal.
6. A process according to any of the above claims, wherein recursively and non-recursively filtered signals are input to non-linear filters for removing specific components of the signals.
7. A process according to Claim 6, wherein it is the noise component of a given signal which is filtered.
8. A process according to any of the Claims 2 to 7, wherein recursively and non-recursively filtered signals are input to band-splitting

filters, and sub-bands of the recursively and non-recursively filtered signals are combined.

9. A process according to Claim 8, wherein sub-bands of corresponding frequency from the respective recursively and non-recursively filtered signals are combined.

10. A process according to Claim 8 or Claim 9, wherein the set of sub-bands of the recursively and non-recursively filtered signals are both input to a single array of non-linear filters, the output of which being a signal multiplexed between the filtered recursive signal, and the filtered non-recursive signal.

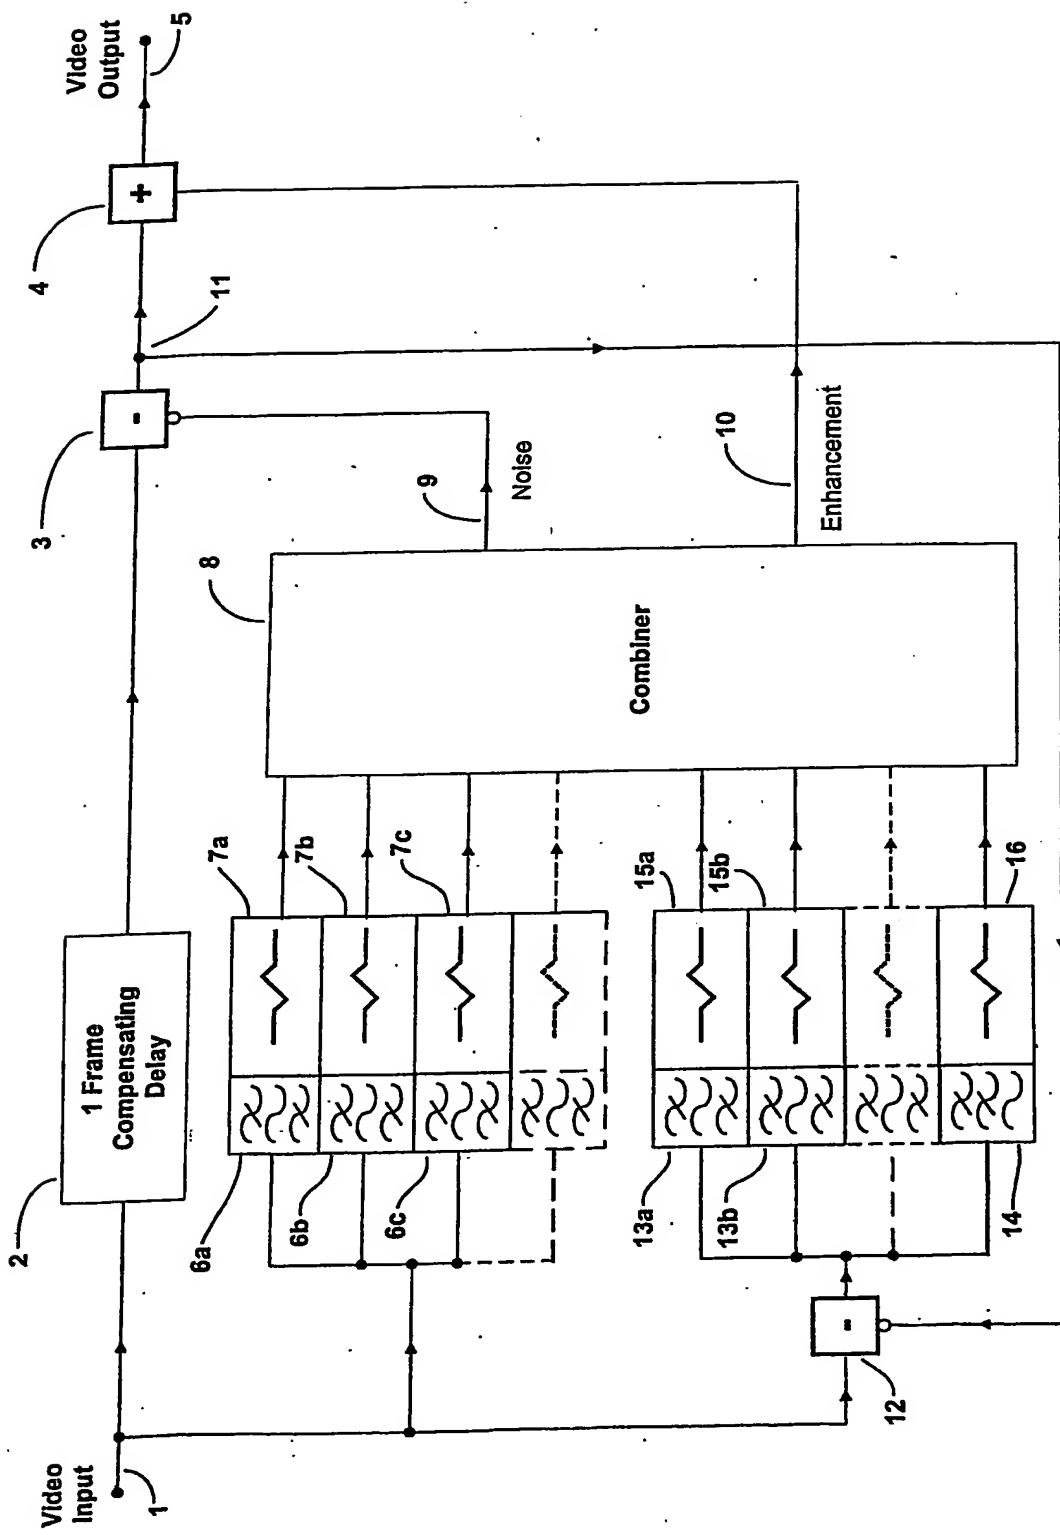


Figure 1: Video Processor

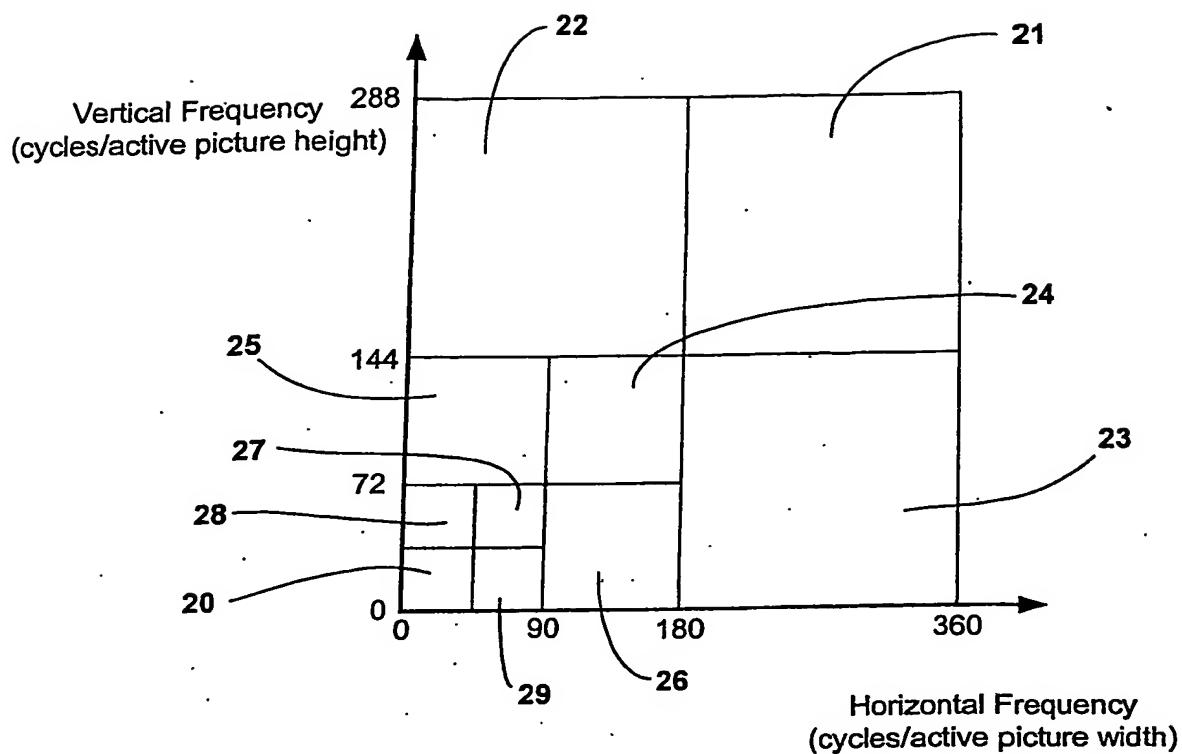


Figure 2: Spatial Frequency Bands

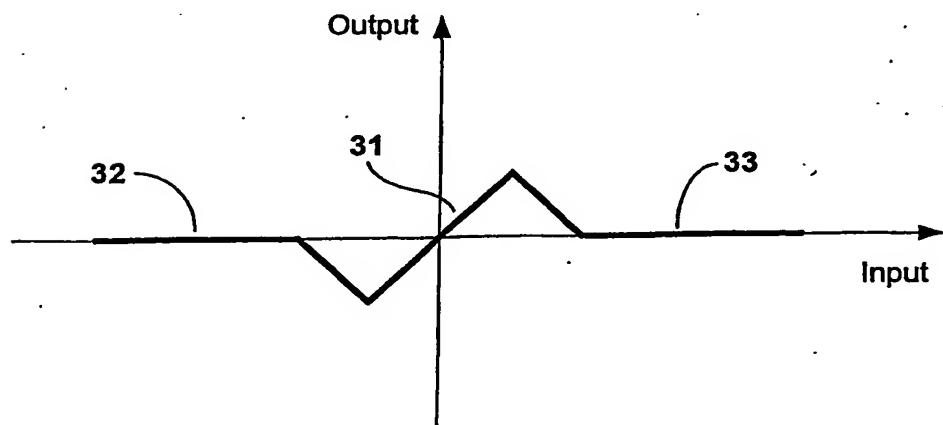


Figure 3: Inverse Coring Characteristic

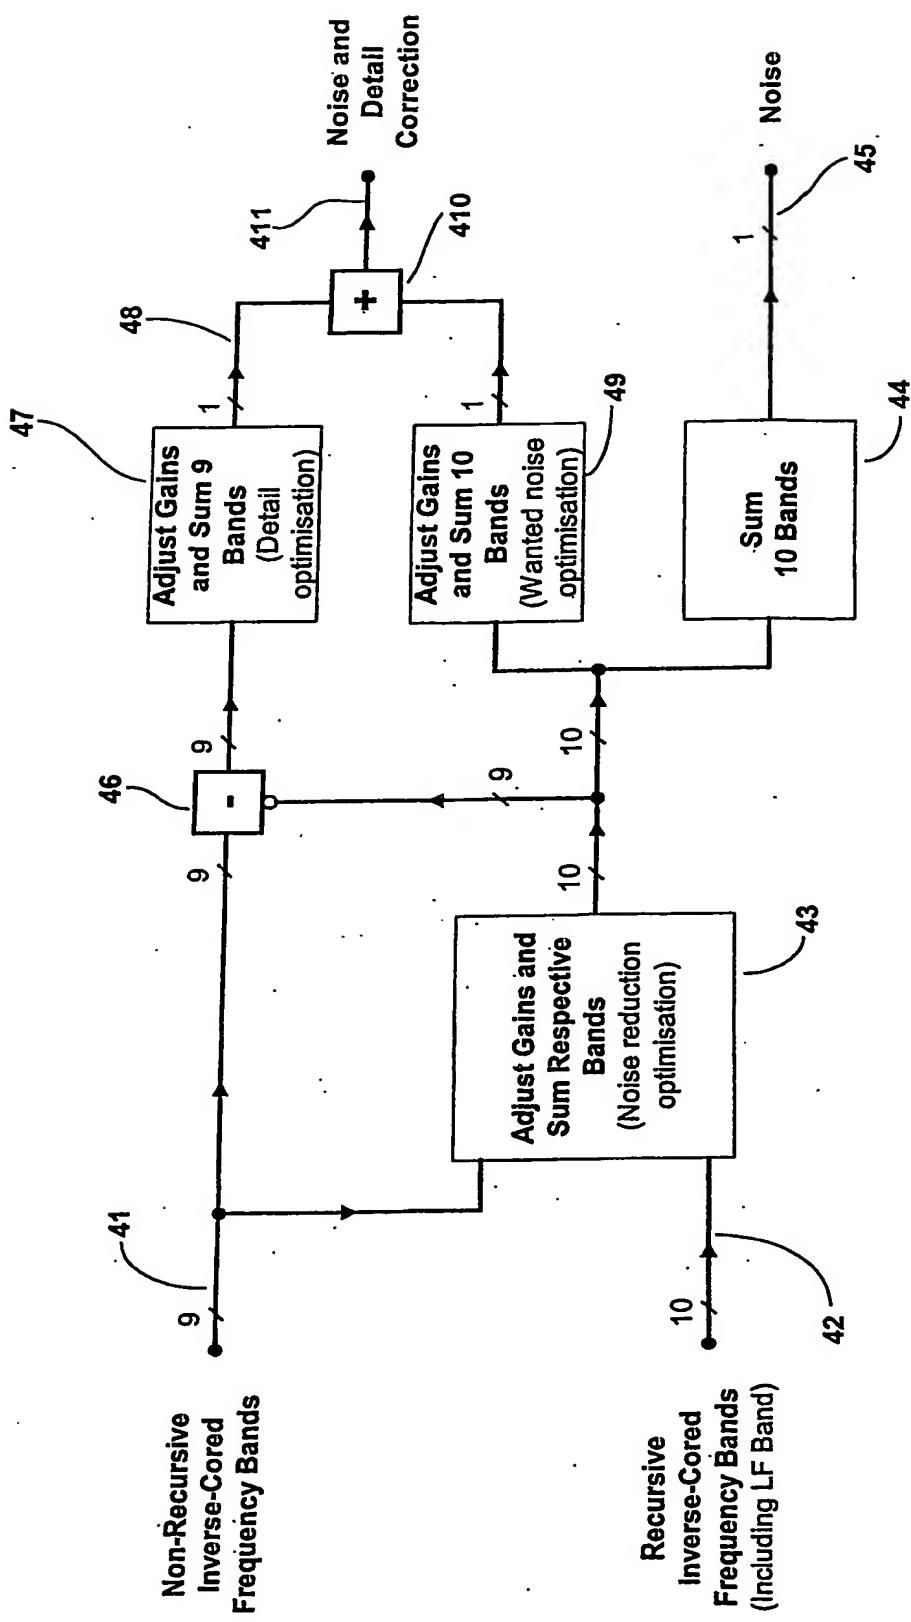


Figure 4: Noise and Detail Processing

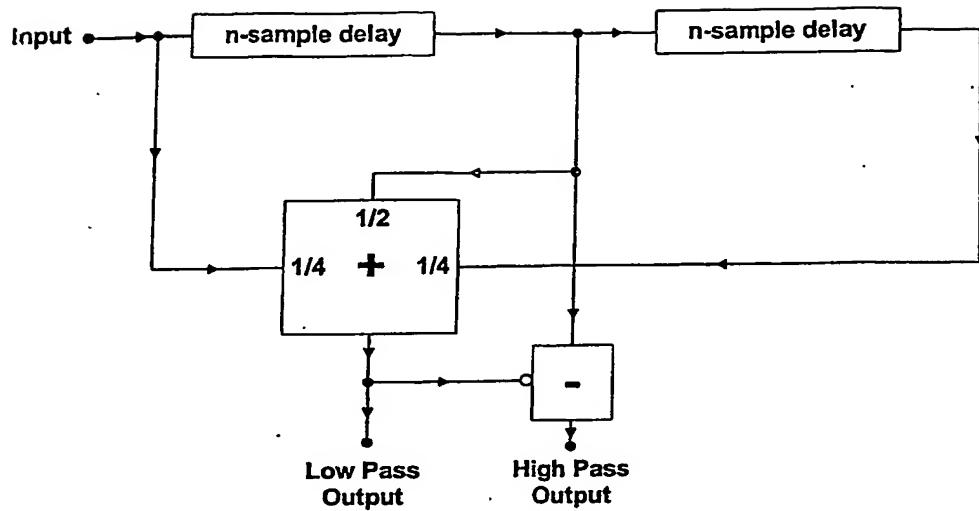


Figure 5a: Band Splitting Filter

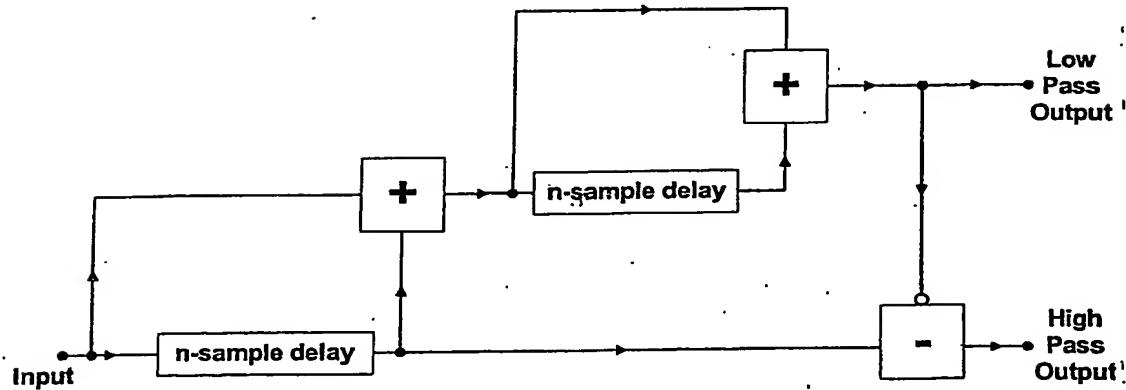


Figure 5b: Band Splitting Filter (Alternative Structure)

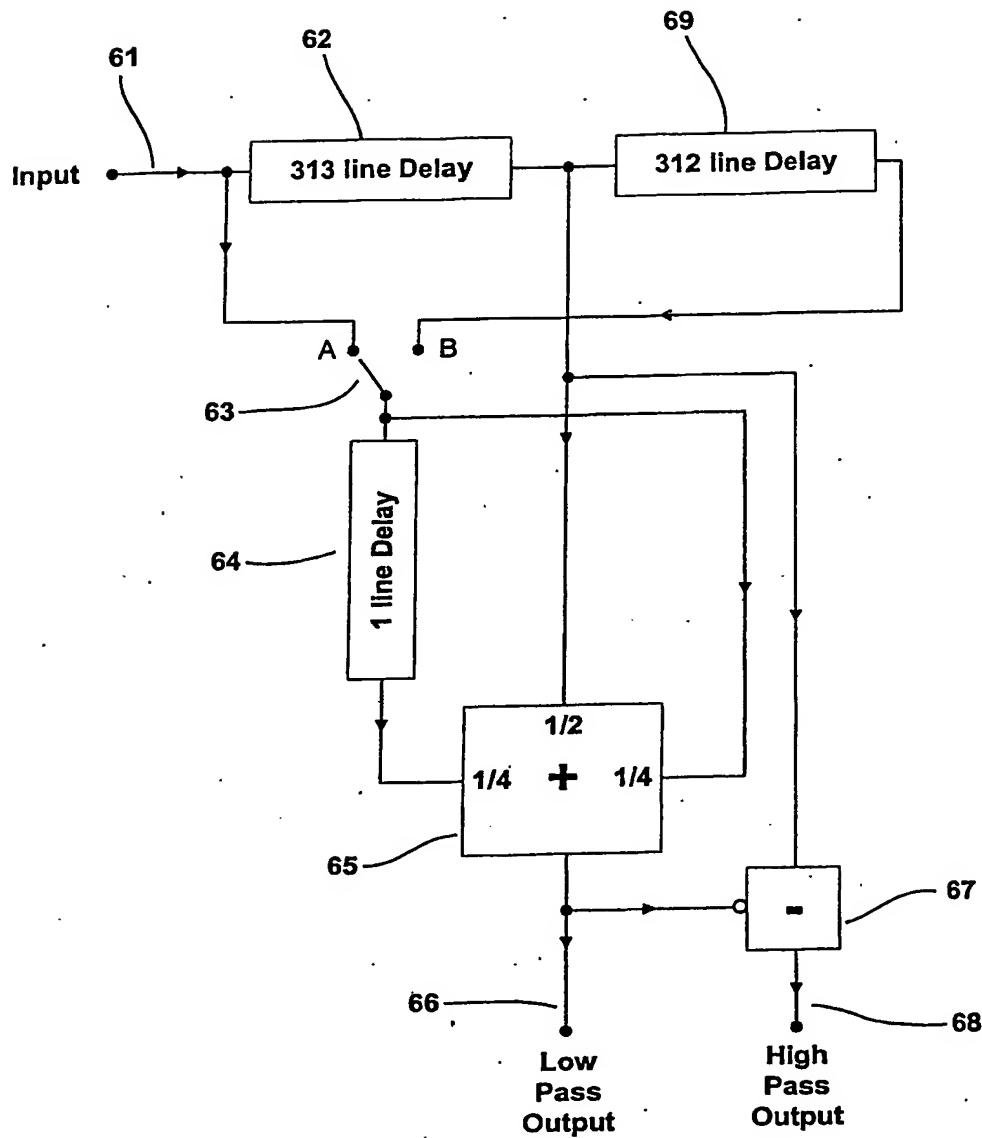


Figure 6: Vertical Band Splitter

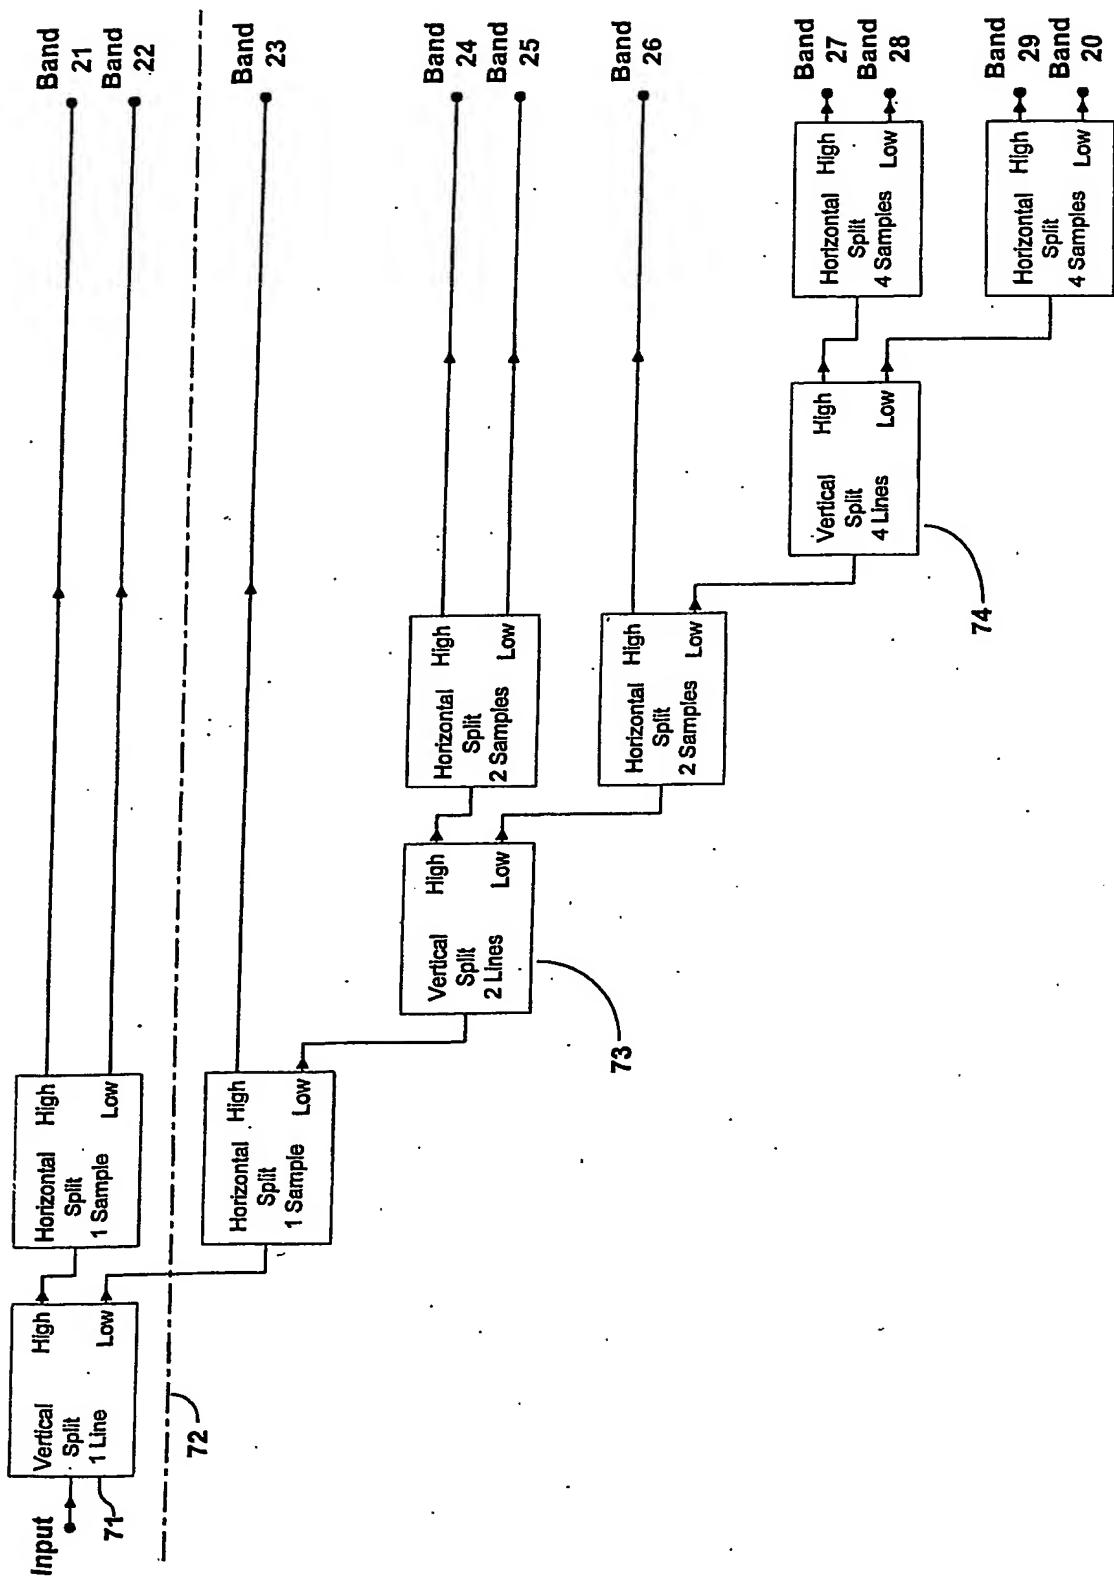


Figure 7: Band Splitting Arrangement

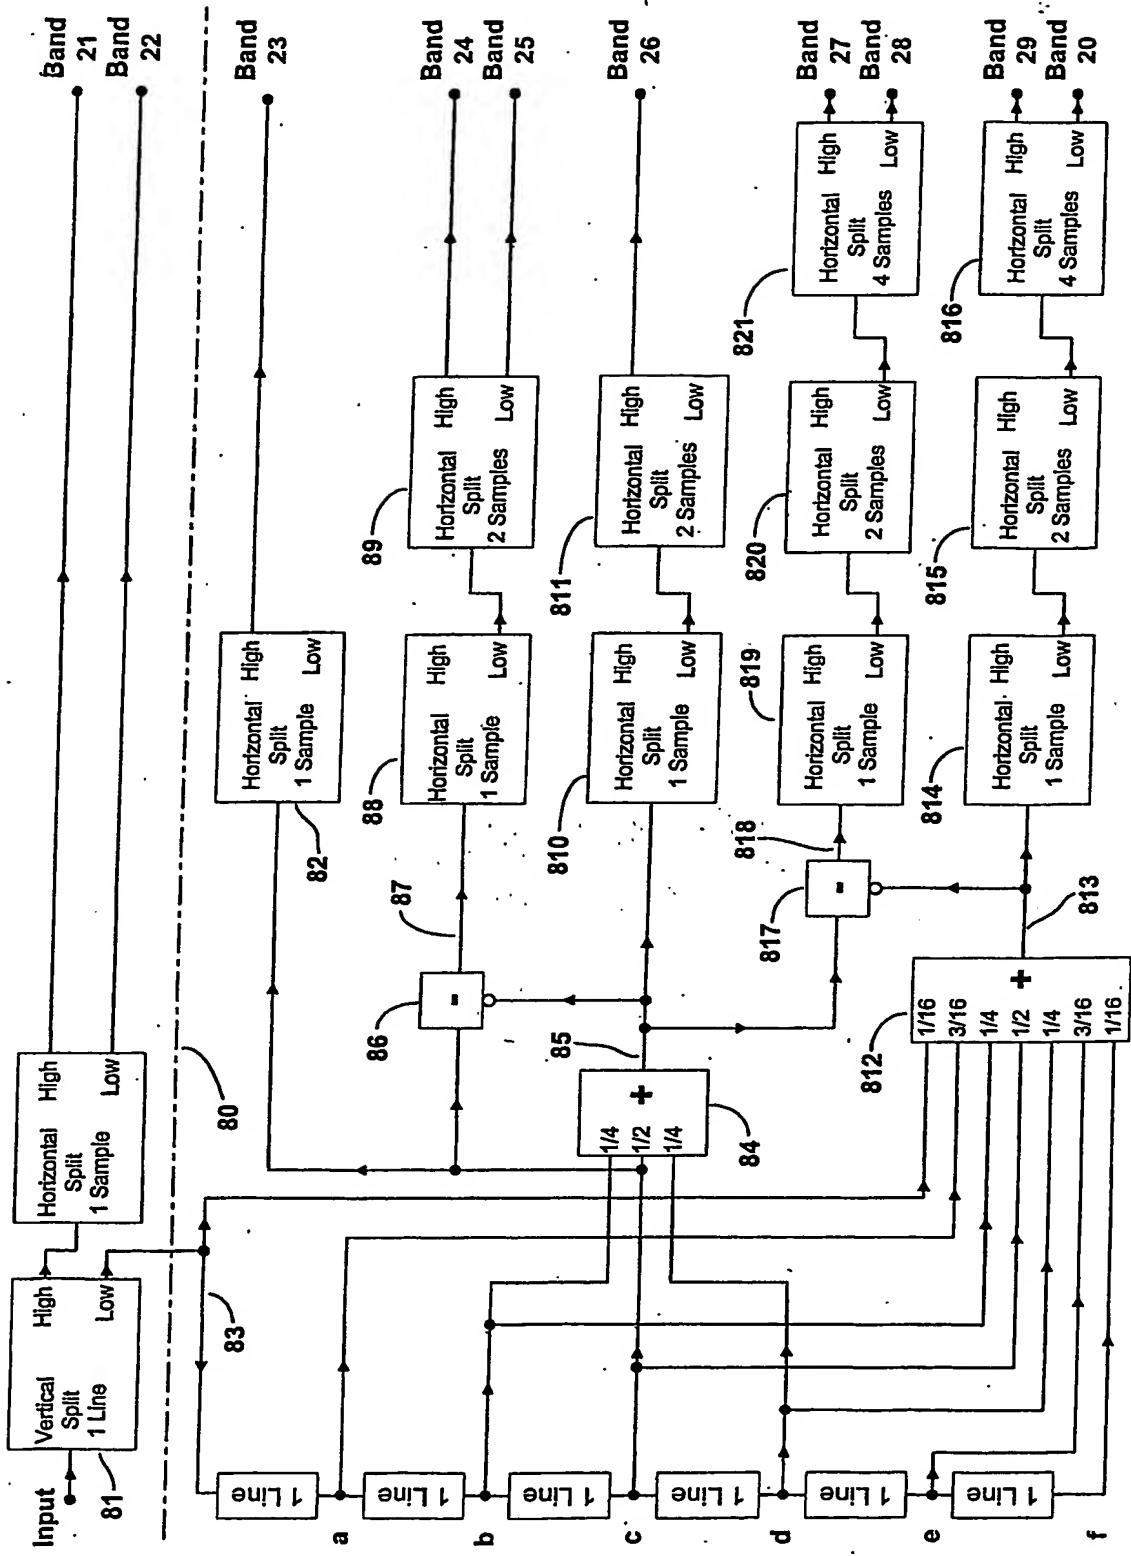


Figure 8: Alternative Band Splitting Arrangement